IN THE UNITED STATES DISTRICT COURT
FOR DISTRICT OF DELAWARE

ORCA SECURITY LTD.,

Plaintiff,

C.A. No. 23-00758-JLH-SRF

v.

WIZ, INC.

Defendants.

WIZ, INC.'S ANSWERING BRIEF IN OPPOSITION TO ORCA SECURITY LTD.'S MOTION TO DISMISS WIZ, INC.'S AMENDED COUNTERCLAIM COUNT IV UNDER FED. R. CIV. P. 12(b)(6)

OF COUNSEL:

Jordan R. Jaffe
Catherine Lacey
Callie Davidson
Alex Miller
WILSON SONSINI GOODRICH & ROSATI, P.C.
One Market Plaza
Spear Tower, Suite 3300
San Francisco, CA 94105
(415) 947-2000

Dated: September 19, 2024

RICHARDS, LAYTON & FINGER, P.A.
Frederick L. Cottrell, III (#2555)
Kelly E. Farnan (#4395)
Christine D. Haynes (#4697)
One Rodney Square
920 N. King Street
Wilmington, DE 19801
(302) 658-6541
cottrell@rlf.com
farnan@rlf.com
haynes@rlf.com

Counsel for Defendant Wiz, Inc.

TABLE OF CONTENTS

TABLE OF AUTHORITIES

Page(s)

CASES

STATUTES

I. NATURE AND STAGE OF PROCEEDINGS

Plaintiff and Counterclaim-Defendant Orca Security Ltd. ("Orca") filed its original complaint initiating this suit on July 12, 2023. D.I. 1. Orca amended its complaint twice and currently alleges that Defendant and Counterclaim-Plaintiff Wiz, Inc. ("Wiz") infringes six patents: U.S. Patent Nos. 11,663,031, 11,663,032, 11,693,685, 11,726,809, 11,740,926, and 11,775,326. D.I. 15. On June 4, 2024, Wiz answered and filed counterclaims for Orca's infringement of U.S. Patent Nos. 11,722,554, 11,929,896, 11,936,693, 12,001,549 (the "'549 patent"), and 12,003,529 (collectively, the "Wiz Asserted Patents"). D.I. 70 at 37-147. On July 25, 2024, Orca moved to dismiss Wiz's Counterclaim IV for Orca's infringement of the '549 patent under Rule 12(b)(6), arguing that the '549 patent lacks patent-eligible subject matter under 35 U.S.C. § 101. D.I. 111. On August 22, 2024, Wiz filed amended counterclaims including additional factual allegations regarding the '549 Patent that addressed any issues raised in Orca's motion to dismiss. *See* D.I. 124, ¶¶ 89-124. Nevertheless, on September 5, 2024 Orca filed the instant Motion to Dismiss, which is largely the same as its prior motion. D.I. 137, 138. Orca has not moved to dismiss Wiz's counterclaims under the other four Wiz Asserted Patents.

II. SUMMARY OF ARGUMENT

1. The asserted claims of Wiz's '549 patent were filed on January 31, 2024, from a parent application originally filed August 28, 2023 and issued on June 4, 2024. These brand-new claims recently approved by the Patent Office provide an improved technological solution that uses cutting-edge generative artificial intelligence ("AI") technology in the context of cloud cybersecurity. This is a quintessential technical problem with no brick-and-mortar analog. Orca's motion to dismiss does not contend otherwise. The claims are directed to an improved system and method for improved cloud cybersecurity incident response, and thus are directed to improvements in the functioning of a computer, rather than an abstract idea. The claims further

require a specific form of generative AI—large language models ("LLMs")—and a specific solution to improve cybersecurity incident response that was neither conventional nor known in the prior art. *See, e.g.*, D.I. 124, ¶¶ 94-97. Indeed, Orca provides no arguments that novel cloud cybersecurity solutions using generative AI could be considered "routine" or "conventional" at all. Orca's attempt to argue that a specific cloud cybersecurity solution using a specific form of generative AI technology should be considered abstract should be rejected.

2. The asserted claims of the '549 patent claim subject matter that is eligible for patenting under 35 U.S.C. § 101. Under the two-step test set out by the Supreme Court in *Alice Corp Pty. v. CLS Bank International*, patent claims are eligible if they (i) are not directed to a patent-ineligible concept, or if (ii) their claims, either individually or as an ordered combination, add an inventive concept to the patent-ineligible concept. 573 U.S. 208, 217 (2014).

3. At *Alice* Step One, the claims are not directed to an abstract idea or any patent-ineligible concept. This ends the inquiry. The independent claims are directed to a specific, technical solution to a problem in the field of computing, including utilizing LLMs in a specific way to address the prior art systems' deficiencies in leveraging inputs and queries in human-understandable natural language in systems also utilizing structured data solutions that are extremely useful for computer systems. The Federal Circuit has long held such claims "necessarily rooted in computer technology in order to overcome a problem specifically arising in the realm of computer networks" are not directed to abstract ideas. *DDR Holdings, LLC v. Hotels.com, L.P.*, 773 F.3d 1245, 1257 (Fed. Cir. 2014); *see also Enfish, LLC v. Microsoft Corp.*, 822 F.3d 1327, 1339 (Fed. Cir. 2016). The claims further improve cybersecurity responses, "improv[ing] the efficient functioning of computers," and are not directed to a patent-ineligible idea. *Data Engine Techs. LLC v. Google LLC*, 906 F.3d 999, 1009 (Fed. Cir. 2018).

4. The inquiry should end at *Alice* Step One, but if not, the elements of the claims, taken individually or as an ordered combination, add an inventive concept that render the claims patent eligible. *See, e.g.*, D.I. 124, ¶¶ 94-108. The specific application of LLMs to cloud cybersecurity incident response are not "well-understood, routine, [and] conventional" as specifically alleged by Wiz in response to Orca's original motion to dismiss. *Aatrix Software, Inc. v. Green Shades Software, Inc.*, 882 F.3d 1121, 1128 (Fed. Cir. 2018) (modification in original). Orca's only argument to the contrary relies on contradicting the well-pled allegations in Wiz's counterclaims, which is improper on a motion to dismiss. *Id.*

5. Wiz's amended counterclaims were directly responsive to Orca's original motion to dismiss. They include specific, plausible factual allegations that the claim elements of the '549 patent were not well-known, routine, or conventional, whether individually or as an ordered-combination. D.I. 124, ¶¶ 95, 103-106. This "precludes dismissal" under 35 U.S.C. § 101 at the Rule 12(b)(6) stage because the Court is required "to resolve any plausibly alleged factual issues in favor of the patentee." *Blackbird Tech v. Uber Techs., Inc.*, C.A. No. 19-561 (MN), 2020 WL 58535, at *6 (D. Del. Jan. 6, 2020).

III. STATEMENT OF FACTS

A. The '549 Patent Specification

The '549 patent specification describes specific techniques for addressing a problem unique to the field of computers and for responding to cybersecurity incidents, as Wiz has explained in its detailed, factual allegations in the amended counterclaims in filed after Orca's original motion to dismiss. *See* D.I. 124, ¶¶ 95-105.

The '549 patent explains that "[s]tructured data solutions are extremely useful for computer systems" because "a data structure, such as a SQL database, makes it easier for a machine to store data, retrieve data, manage data." '549 patent at 1:31-35. On the other hand, in the context of cybersecurity solutions, queries and alerts when presented in natural language form can lack context and important information, even using prior art natural language processing techniques. *Id.* at 1:38-50; *see also* D.I. 124, ¶ 94.

The ’549 patent, and in particular the asserted claims, provide a specific, technical solution to this deficiency in the prior art using cutting-edge artificial intelligence technology that leverages both natural language and structured data techniques in an unconventional manner. *See, e.g.*, D.I. 124, ¶ 95. The claims are directed to a specific improvement in computer capabilities, using not just artificial intelligence (“AI”) in general, but a specific application of AI—large language models, or “LLMs”—to improve cybersecurity incident response, security databases and mitigation actions. *Id.*; *see also* ’549 patent at 10:27-28 (“In some embodiments, the [artificial neural network] is a large language model, such as GPT, BERT, and the like.”). This improved technique and solution for cybersecurity incident response includes generating a prompt for an LLM based on an incident input based on a cybersecurity event, where the LLM generates an output based on the generated prompt, mapping the received incident input into a scenario of a plurality of scenarios associated with an incident response based on the output of the LLM, generating a query based on both the received incident input and the mapped scenario, executing a query on a security database that includes representation of a computing environment, and initiating a mitigation action based on a result of the executed query. *See* ’549 patent at 4:23-41; *see also* D.I. 124, ¶ 96.

Contrary to Orca’s contention that the advance is merely “supplying additional context or ‘structure’ for natural language queries,” the ’549 patent techniques involve generating a specific prompt for an LLM where the LLM maps an incident input into a scenario and generates a query for the database based on both the received incident input and the mapped scenario; running the query on the database allows initiation of a mitigation action based on the result of the executed query. *See, e.g., id.*, ¶ 97. This is an improved cloud cybersecurity system, rather than one merely supplying "additional context." These techniques were not well-understood, routine, or conventional. *See*. D.I. 124, ¶¶ 98-101.

B. The Asserted '549 Patent Claims

1. The Independent Claims

Wiz asserts claims 1-5 and 11-16 of the '549 patent against Orca. Claim 1 recites:

1. A method for providing cybersecurity incident response, comprising:

receiving an incident input based on a cybersecurity event;

generating a prompt for a large language model (LLM) based on the received incident input;

configuring the LLM to generate an output based on the generated prompt;

mapping the received incident input into a scenario of a plurality of scenarios based on the output of the LLM, wherein each scenario is associated with an incidence response;

generating a query based on the received incident input and the mapped scenario;

executing the query on a security database, the security database including a representation of a computing environment; and

initiating a mitigation action based on a result of the executed query.

Claim 11 recites: "A non-transitory computer-readable medium storing a set of instructions for providing cybersecurity incident response, the set of instructions comprising: one or more instructions that, when executed by one or more processors of a device, cause the device to" perform essentially the same method steps as claim 1. Claim 12 recites: "A system for providing cybersecurity incident response comprising: a processing circuitry; a memory, the memory containing instructions that, when executed by the processing circuitry, configure the system to" similarly to perform the method of claim 1.

2. The Asserted Dependent Claims

As discussed below, the asserted dependent claims add meaningful limitations to the independent claims. Claims 3, 5, 14, and 16 include specific limitations as to how to train the LLM utilized in the solution of the independent claims. *See* claims 3 and 14 ("the LLM is trained on any one of: a data schema utilized in representing the computing environment, incident data classified to a scenario, the plurality of scenarios, and a combination thereof"); claims 5 and 16 ("training the LLM further on a plurality of database queries, each database query executable on the security database"). Claims 4 and 15 add the limitation of generating a second prompt based on specific elements: "generating a second prompt for the LLM which when executed by the LLM outputs the query, wherein the second prompt is generated based on any one of: the received incident input, the data schema, the plurality of scenarios, and a combination thereof." Claims 2 and 13 further narrow the incident inputs.

IV. LEGAL STANDARD

A. Motions to Dismiss Under Rule 12(b)(6)

Patent eligibility can be resolved at the Rule 12(b)(6) stage "only when there are no factual allegations that, taken as true, prevent resolving the eligibility question as a matter of law." *Aatrix Software*, 882 F.3d at 1125. The underlying question in the Section 101 analysis "of whether a claim element or combination of elements is well-understood, routine and conventional to a skilled artisan in the relevant field is a question of fact." *Berkheimer v. HP Inc.*, 881 F.3d 1360, 1368 (Fed. Cir. 2018). Thus, where "plausible factual allegations that the claimed invention improves upon the prior conventional systems" raise a dispute regarding "whether the claim elements and their ordered combination is simply well-known, routine and conventional," such a dispute precludes dismissal. *Blackbird Tech*, 2020 WL 58535 at *6.

B. Patent Eligibility under 35 U.S.C. § 101

The Supreme Court has established a two-step test for determining whether patent claims are invalid under § 101. *Alice Corp.*, 573 U.S. at 217 (2014). In Step One, the Court must "determine whether the claims at issue are directed to a patent-ineligible concept." *Id.* If the answer is no, the inquiry ends, and the claims are patent eligible. *See, e.g., Finjan, Inc. v. Blue Coat Sys., Inc.*, 879 F.3d 1299, 1305-06 (Fed. Cir. 2018). Otherwise, at Step Two, the Court "consider[s] the elements of each claim both individually and as an ordered combination" to determine if there is an "inventive concept— *i.e.*, an element or combination of elements that is sufficient to ensure that the patent in practice amounts to significantly more than a patent upon the [ineligible concept] itself." *Alice*, 573 U.S. at 217-18 (internal quotations and citations omitted). "Claims pass muster at step two when they 'involve more than performance of well-understood, routine, and conventional activities previously known to the industry.'" *Trackthings LLC v. Netgear, Inc.*, C.A. No. 22-981-RGA, 2023 WL 4926184, at *10 (D. Del. Aug. 2, 2023), *report and recommendation adopted*, 2023 WL 5993186 (D. Del. Sept. 15, 2023). Issued claims are presumed patentable because "the Patent and Trademark Office has already examined whether the patent satisfies 'the prerequisites for issuance of a patent,' including § 101." *Cellspin Soft, Inc. v. Fitbit, Inc.*, 927 F.3d 1306, 1319 (Fed. Cir. 2019).

V. ARGUMENT

The asserted claims of the '549 patent are eligible under Section 101 and *Alice* because they are not directed to ineligible subject matter. They also contain an inventive concept, and in any event, fact issues preclude resolving this issue on summary judgment let alone the pleading stage.

A. *Alice* Step One: The Claims Are Not Directed to an Abstract Idea

1. The Independent Claims Are Directed to A Specific Solution to a Computer Problem

The asserted ’549 patents claims are not directed to an abstract idea, or any other patent-ineligible subject matter. The inquiry thus ends at *Alice* Step 1.

First, the claims are not directed to an abstract idea because they are “necessarily rooted in computer technology in order to solve a specific problem in the realm of computer networks.” *Packet Intel. LLC v. NetScout Sys., Inc.*, 965 F.3d 1299, 1309 (Fed. Cir. 2020); *see also DDR Holdings*, 773 F.3d at 1257. As explained in the ’549 patent specification and Wiz’s counterclaims, specifically in the context of cybersecurity solutions, queries and alerts when presented in natural language form can lack context and important information, such as the relevant workloads, root causes, or potential mitigation whereas computers typically communicate using structured data, such as SQL for databases. D.I. 124, ¶ 94; *see also* 1:22-43; 7:30-34. Prior incident response cybersecurity systems failed to provide a solution that improved a cybersecurity incident response system while leveraging both natural language processing and structured data. D.I. 124, ¶ 94; *see also* ’549 patent at 1:44-56.

The asserted claims address this problem arising in the realm of computer networks with a solution entirely rooted in computer technology. *See* D.I. 124, ¶¶ 94, 97. The asserted claims utilize one form of generative AI—LLMs—in a specific way to address this problem. For example, the claims require using generative AI to map what specific cloud cybersecurity response scenarios are applicable given the particular received incident input. *See* ’549 patent claim 1; D.I. 124, ¶¶ 95-97. The claims, therefore, “do not merely recite the performance of some business practice known from the pre-Internet world along with the requirement to perform it on the Internet," but instead are directed to an improved cloud cybersecurity system that utilizes generative AI to improve incident response. *DDR Holdings*, 773 F.3d at 1257.

Moreover, the claims are "directed to a ***specific implementation*** of a solution to a problem in" cybersecurity systems and are therefore non-abstract. *Enfish*, 822 F.3d at 1339 (emphasis added). The claims recite ***how*** the claimed method, system, and computer readable medium are implemented. The claims recite generating a prompt based on an incident input based on a cybersecurity event, mapping that received incident input into a scenario of a plurality of scenarios based on the output of the LLM, wherein each scenario is associated with an incident response, followed by specific further steps on how the improved system or method uses the output from the artificial intelligence model to generate a query for the security database based on the received incident input and the mapped scenario, to execute such a query, and then to initiate a mitigation action. *See* D.I. 124, ¶ 100; '549 patent at 4:23-42, claim 1. The claims do not recite "using AI" with no details, but rather provide a specific use of a specific type of AI—an LLM—and how that LLM is used, including mapping the received incident input into a scenario of a plurality of scenarios based on the output of the LLM, wherein each scenario is associated with an incident response, followed by specific further steps on generating and executing a query for the security database, and then initiating a mitigation action. *See* D.I. 124, ¶ 100; '549 patent at 5:16-34, claim 1.

The claims further require that the database be a "security database" that "includ[es] a representation of a computing environment" rather than any generic database. *See* D.I. 124, ¶ 101; claim 1. In the claimed solution, the scenarios are mapped to one of a plurality of scenarios based on an incident response, which can then be used directly to interface with the security database that contains the representation of the computing environment. The asserted claims are thus directed to a particular enhanced cybersecurity system that requires, for example, mapping what cloud cybersecurity response scenarios are applicable given received incident input utilizing an LLM. *See* ’549 patent claim 1; D.I. 124, ¶¶ 95-97. *Id.* Claims directed to a specific implementation to achieve improvements in *inter alia* querying databases are not directed to an abstract idea. *See Enfish*, 822 F.3d at 1333 (claims directed to specific implementation for “faster searching of data” in databases was non-abstract).

The claims are also non-abstract because they are directed to “improv[ing] the efficient functioning of computers,” and are thus not directed to a patent-ineligible idea. *Data Engine*, 906 F.3d at 1009. The claims are directed to a solution that improves computer capabilities that allows improved responses to cybersecurity incidents. *See* D.I. 124, ¶¶ 95- 101. The claims of the ’549 patent are directed to a specific, technical improvements that “increase[] the usability of a cybersecurity monitoring solution, and improves the incident response time” and “utiliz[e] a large language model to map an incident input to a scenario [that]. . . decreases incidence response time, and therefore decreases time to mitigation in the event of a cybersecurity breach.” ’549 patent at 8:31-38; *see also* D.I. 124, ¶ 99. The ’549 patent’s “claimed advance is a concrete assignment of specified functions among a computer’s components to improve computer security, and this claimed improvement in computer functionality is eligible for patenting.” *Ancora Techs., Inc. v. HTC Am., Inc.*, 908 F.3d 1343, 1344 (Fed. Cir. 2018), as amended (Nov. 20, 2018); *Finjan*, 879 F.3d at 1305 (“The asserted claims are therefore directed to a non-abstract improvement in computer functionality, rather than the abstract idea of computer security writ large.”)

The claims of the ’549 patent are directed to specific, novel implementations of computer technology to solve a specific problem rooted in cybersecurity systems, and thus are non-abstract. Orca’s arguments to the contrary are not compelling, and broadly rely on improperly “overgeneralizing claims.” *TecSec, Inc. v. Adobe Inc.*, 978 F.3d 1278, 1293 (Fed. Cir. 2020).

When Orca asserts the "claims focus on retrieving, contextualizing, querying, and responding to cybersecurity threat information" (D.I. 138 at 8), it "characteriz[es] the claims at 'a high level of abstraction' that is 'untethered from the language of the claims'" in an attempt to ensure "the exceptions to § 101 swallow the rule.'" *TecSec*, 978 F.3d at 1293; *see also CardioNet, LLC v. InfoBionic, Inc.*, 955 F.3d 1358, 1371 (Fed. Cir. 2020). This characterization ignores the specific requirements of the claims, including that they require generating a prompt for an LLM based on an incident input based on a cybersecurity event, mapping that incident input to a scenario based on the output of the LLM, generating a query for a security database based on the incident input and scenario, and that the cybersecurity database include a representation of a computing environment.

Orca's three enumerated arguments also miss the mark. First, Orca conflates the idea that the claimed solution may involve human interaction with the idea that it is merely "directed to a *human* problem." D.I. 138 at 8 (emphasis in original). On the contrary, the claims are directed to a problem arising in the realm of cybersecurity computer networks, and provide a solution rooted in computer technology. *See* D.I. 99. The claimed solution uses an LLM to map an incident response to a plurality of cybersecurity scenarios, then the system can formulate a structured data query for the cybersecurity database based on the received incident input and the mapped scenario, execute the query on the database, and execute a mitigation action. *See* D.I. 97. The claims are not directed to abstract ideas merely because they may involve some human interaction. *See, e.g.*, *Core Wireless Licensing S.A.R.L. v. LG Elecs., Inc.*, 880 F.3d 1356, 1360 (Fed. Cir. 2018) (claims requiring "list being selectable" not directed to abstract idea); *CoolTVNetwork.com, Inc. v. Facebook, Inc.*, C.A. No. 19-292-LPS-JLH, 2019 WL 4415283, at *11 (D. Del. Sept. 16, 2019) (claim "requir[ing] user interaction" not directed to abstract idea at Rule 12(b)(6) stage).

Even if the claims were directed to activities theoretically performable by humans, the Federal Circuit has rejected the idea that computing solutions for tasks that are human-performable are abstract where, as here, they involve a specific implementation to solve a problem arising in computing; "processes that automate tasks that humans are capable of performing are patent eligible if properly claimed." *McRO, Inc. v. Bandai Namco Games Am., Inc.*, 837 F.3d 1299, 1313 (Fed. Cir. 2016). None of the cases cited by Orca speak to the contrary. *Trinity Info Media, LLC v. Covalent, Inc.*, 72 F.4th 1355, 1363 (Fed. Cir. 2023) (claims directed to "how to improve existing polling systems by performing progressive polling, not how to improve computer technology."); *PersonalWeb Techs. LLC v. Google LLC*, 8 F.4th 1310, 1317 (Fed. Cir. 2021) (claims directed to "using a content-based identifier to perform an abstract data-management function"); *Intell. Ventures I LLC v. Symantec Corp.*, 838 F.3d 1307, 1313 (Fed. Cir. 2016) (claims directed to generic idea of "filtering files/e-mail"); *Int'l Bus. Machines Corp. v. Zillow Grp., Inc.*, 50 F.4th 1371, 1381 (Fed. Cir. 2022) (claims directed to display of information); *Chewy, Inc. v. Int'l Bus. Machs. Corp.*, 94 F.4th 1354, 1365 (Fed. Cir. 2024) (patentee admitted claims were directed only to "obtaining search results from a user's search query and using those search results to identify targeted advertisements"); *Trading Techs. Int'l, Inc. v. IBG LLC*, 921 F.3d 1084, 1089, 1092 (Fed. Cir. 2019) (claims merely directed to "a financial trading method used by a computer" or "the abstract idea of graphing (or displaying) bids and offers to assist a trader to make an order").

In *Simio, LLC v. FlexSim Software Products*, the patent specification made clear that its "key advance" was "using graphics instead of programming to create object-oriented simulations" but also acknowledged that "using graphical processes to simplify simulation building ha[d] been done since the 1980s and 1990s." 983 F.3d 1353, 1359-60 (Fed. Cir. 2020). Therefore, the claimed advance was simply "applying the already-widespread practice of using graphics instead of programming to the environment of *object-oriented* simulations [and] no more than an abstract idea." *Id.* Orca has not and cannot point to any similar record here that the techniques claimed in the '549 patent were conventional and merely being applied in a new area. There is nothing in the specification or elsewhere in the record suggesting that generating a prompt for an LLM based on an incident input, mapping the input to one of many scenarios based on an output from the LLM, and using the input and the scenario to generate a database query was a widespread practice that the '549 patent claims merely applied to the cybersecurity environment. In fact, the specification indicates the opposite. *See* '549 patent at 1:44-45 (noting that problem addressed by the patent is a "recurring issue" with prior art natural language processing techniques); *see also* D.I. 124, ¶¶ 96-98. None of the cases cited by Orca that relate to simply computerizing human activity using generic techniques compare to the claims of the '549 patent, which claim specific, novel techniques to solve a problem unique to computers.

Orca's second argument that the claims do not provide any specific improvement in computer technology fares no better, for much the same reasons as discussed above. D.I. 138 at 10. The claimed solution provides a specific improvement in computer technology for leveraging both natural language processing and structured data to improve querying a security database and cybersecurity response. *See, e.g.*, D.I. 124, ¶¶ 94-101. Orca's reliance on *FairWarning IP, LLC v. Iatric Systems, Inc.*, 839 F.3d 1089 (Fed. Cir. 2016) is misplaced. D.I. 138 at 11. In that case, the claims "merely implement[ed] an old practice in a new environment" because they simply computerized a set of rules that asked "the same questions . . . that humans in analogous situations detecting fraud have asked for decades, if not centuries." *FairWarning*, 839 F.3d at 1094-95. Orca provides no explanation of what it means when it asserts that the claimed steps of the '549 patent are merely steps humans have performed for decades. *See* D.I. 138 at 11. Orca identifies no long-standing human activity analogous to, for example, utilizing an LLM for mapping an incident input based on a cybersecurity event into one of a plurality of scenarios associated with an incidence response, and generating a query for a database based on the received incident input and the mapped scenario. *Id.* In any event, Orca's assertion is contradicted by the plausible, factual allegations of Wiz's counterclaims, and therefore must be disregarded at this stage. *See, e.g.*, D.I. 124, ¶ 99.

BSG Tech LLC v. Buyseasons, Inc., 899 F.3d 1281 (Fed. Cir. 2018) is no help to Orca, either. The claims in that case recited broadly providing users with historical usage information of parameters to encourage them to "maintain consistency in how different users describe items;" the claims at issue did not claim specific technical solutions to providing that information, and thus amounted to claims on "having users consider previous item descriptions before they describe items to achieve more consistent item descriptions," which does not solve a problem specific to computers. *Id.* at 1284, 1286. By contrast, the '549 patent claims relate to a specific technical solution for generating a query of a security database, and executing mitigation actions, to respond to a cybersecurity incident leveraging natural language and structured data.

Orca's third argument is also wrong. The '549 patent claims do not merely "recite basic computerized steps for collecting, analyzing, and presenting information" "without any explanation of *how* the system accomplishes its goals." D.I. 138 at 13 (emphasis in original). The opposite is true. The claims teach a specific solution for cybersecurity response involving the techniques of generating a prompt for an LLM, mapping an incident input to a scenario based on the output of the LLM, and generating the query based on both the incident input and the scenario, and executing the query and mitigation. *See* claim 1. Nor are the claims analogous to those found ineligible in *Universal Secure Registry LLC v. Apple Inc.*, 10 F.4th 1342 (Fed. Cir. 2021) ("*USR*"). In *USR*, the claim included no limitation that the mapping involve an LLM, involve mapping to an incident input into a scenario associated with an incidence response, or on how the “mapping” was performed at all—the limitation was merely to map a “code to an identity.” *Id.* at 1348. The ’549 claims are limited to the technique of generating a prompt for an LLM, the LLM generating an output based on the prompt, and mapping the received incident input into a scenario of a plurality of scenarios, wherein each scenario is associated with an incidence response. *See* ’549 patent at claim 1.

2. The Asserted Dependent Claims Are Directed to Even More Concrete Subject Matter

The dependent claims include further elements that confirm they relate to specific, technical solutions in the field of computer security, which Orca fails to address. Orca asserts only in conclusory fashion that claim 1 is representative, but the dependent claims include additional limitations pertinent to the Section 101 analysis. *Berkheimer*, 881 F.3d at 1365 (“A claim is not representative simply because it is an independent claim.”).

The additional limitations of claims 3, 5, 14, 16 require the LLM utilized in the solution claimed in the independent claims be trained on specific types of data. *See* ’549 patent at claim 3 (“the LLM is trained on any one of: a data schema utilized in representing the computing environment, incident data classified to a scenario, the plurality of scenarios, and a combination thereof”); claim 14 (same); claim 5 (“training the LLM further on a plurality of database queries, each database query executable on the security database”); claim 16 (same). For example, the specification explains that the LLM may “include[] a fine-tuning mechanism” which “allows [*sic*] to freeze some weights of a neural network while adapting others based on training data which is unique to a particular set of data.” *Id.* at 10:41-45. These claims further specify the implementation of the claimed solution.

Claims 4 and 15 add the limitation of generating a second prompt based on specific elements: “generating a second prompt for the LLM which when executed by the LLM outputs the query, wherein the second prompt is generated based on any one of: the received incident input, the data schema, the plurality of scenarios, and a combination thereof." This is another concrete, technical step in the solution that Orca fails to meaningfully address. Claims 2 and 13 further limit the claimed incident input to "any one of: a query, a statement, and a combination thereof." These claims are not limited to an abstract idea at least because they depend from the independent claims, which are not directed to an abstract idea as discussed above.

As it did for the independent claims, Orca improperly overgeneralizes the dependent claims as its sole argument as to the asserted dependent claims. *See* D.I. 138 at 14-15.

B. *Alice* Step Two: Inventive Concepts Preclude Dismissal

1. The Independent Claims Include Inventive Concepts

Orca's Motion should be denied at Step One. However, should the Court reach Step Two, Orca falls painfully short of its burden to show that the additional elements of the claims beyond the purported abstract idea were "well-understood, routine, and conventional" to justify dismissal. *Aatrix Software, Inc. v. Green Shades Software, Inc.*, 890 F.3d 1354, 1356 (Fed. Cir. 2018) (Moore, J., concurring in the denial of rehearing *en banc*); *see also Aatrix Software*, 882 F.3d at 1129-30 (overturning dismissal where "nothing in the specification describes [the claimed] importation of data as conventional").

As explained in Wiz's amended counterclaims, the asserted claims and their elements both individually and as an ordered combination are not well-understood, routine, conventional activities. *See, e.g.*, D.I. 124, ¶¶ 98, 103-105. As explained in the '549 patent specification, specifically in the context of cybersecurity solutions, prior incident response cybersecurity systems failed to provide a solution that improved a cybersecurity incident response system while leveraging both natural language processing and structured data. *See, e.g. id.*, ¶ 94. The '549 patent, and in particular the asserted claims, provide a novel improvement directed to a specific improvement in computer capabilities, using not just artificial intelligence ("AI") in general, but a specific application of AI—large language models, or LLMs—to cybersecurity incident response, security databases and mitigation actions. *See, e.g. id.*, ¶ 95. The '549 patent discloses and claims a unique technological solution to this particularly technological problem, providing an improved technique and solution for cybersecurity incident response leveraging LLMs and includes a specific technique to analyze incident input without knowing, for example, which workloads are affected and the particular context of the input. *See, e.g. id.*, ¶ 96. At least several elements, even considered alone, were not well-understood, routine, or conventional activities. *Id.*, ¶¶ 103-104. For example, the use of a particular type of artificial intelligence technology, LLMs, to "map[] the received incident input into a scenario of a plurality of scenarios based on the output of the LLM, wherein each scenario is associated with an incidence response," was neither routine nor conventional in the industry, including as shown by the prosecution history of the '549 patent. *See, e.g. id.*, ¶ 103. Further, "generating a query based on the received incident input to and the mapped scenario," "executing the query on a security database, the security database including a representation of a computing environment," and "initiating a mitigation action based on a result of the executed query" are each individually non-routine and unconventional techniques for cybersecurity systems at the time of filing. *See, e.g.*, *id.*, ¶ 104. These steps more than render the claims patent-eligible. *See CosmoKey Sols. GmbH & Co. KG v. Duo Sec. LLC*, 15 F.4th 1091, 1098 (Fed. Cir. 2021) (finding claims patent-eligible because nothing suggested "the last four claim steps" were routine or conventional).

The claims are also neither routine nor conventional as an ordered combination as well. *See* D.I. 105. This is an independent reason to deny this Motion, as at *Alice* Step Two "an inventive concept can be found in the non-conventional and non-generic arrangement of known, conventional pieces." *Bascom Glob. Internet Servs., Inc. v. AT&T Mobility LLC*, 827 F.3d 1341, 1350 (Fed. Cir. 2016). Nothing in the specification or other record for this Motion indicates the ordered combination of the elements of the claims was well-understood, routine, or conventional, which include generating a prompt for an LLM based on the received incident input, configuring the LLM to generate an output based on the generated prompt, mapping the received incident input into a scenario of a plurality of scenarios based on the output of the LLM, generating a query based on the received incident input and the mapped scenario, executing the query on a security database and initiating a mitigation action based on a result of the executed query. *See, e.g.*'549 patent at claim 1. Indeed, the specification indicates the opposite, as it notes challenges in prior art solutions and that it would "be advantageous to provide a solution that would overcome the challenges." '549 patent at 1:44-58. The '549 patent also issued over numerous prior art references considered during its prosecution. *See id.* at (56).

Orca fails to provide any analysis to support its assertion that there is no inventive concept to the claim elements as an ordered combination. D.I. 138 at 17-18. It simply asserts as much in conclusory fashion and quotes from prior cases without connecting those quotations or cases to ***this case***. This is reason enough to deny the motion at *Alice* Step Two, as the claims must be considered as an "ordered combination."

As to taking the claim elements individually, Orca improperly "describe[es] the claims at such a high level of abstraction and untethered from the language of the claims" as to attempt to ensure the claims appear to contain no inventive concept. *Enfish*, 822 F.3d at 1337. As discussed at length above, the claimed solution specifically requires far more than "using computers to retrieve, contextualize, map, or match data; query database; and perform actions based on results," contrary to Orca's contentions. *See* D.I. 138 at 17. But this abstraction of its own invention is all Orca addresses in its Motion, as well as citations to cases that are inapposite for the reasons discussed at Step One. *Id.* at 16-18. Orca does nothing to show, for example, that mapping a received incident input into a scenario associated with an incidence response based on an output from an LLM, or that using an incident input and a mapped scenario to generate a query of a security database, are "routine, well-understood, or conventional"—let alone to overcome Wiz's plausible, factual, allegations that this is not the case.

Orca also cites a string of cases it contends show the LLM elements of the claims were well-understood, routine, or conventional, but none stand for the proposition that the specific use of an LLM, among other elements, as claimed in the '549 patent is not an inventive concept. D.I. 138 at 17 fn.3. In fact, none of the cases actually address claims that recite LLMs at all, which are a specific type of artificial neural network. *See* '549 patent at 10:26-28 ("In some embodiments, the ANN [artificial neural network] 230 is a large language model, such as GPT, BERT, and the like."); *see also* D.I. 124, ¶ 95. The cases are inapposite in other ways as well. *See Intell. Ventures I LLC v. Cap. One Bank (USA)*, 792 F.3d 1363, 1370-71 (Fed. Cir. 2015) (attorney argument merely referred to "software" "brain"); *Recentive Analytics, Inc. v. Fox Corp.*, 692 F. Supp. 3d 438, 444-45 (D. Del. 2023) (claims directed to generic method of generating "television schedules"); *Hyper Search, LLC v. Facebook, Inc.*, C.A. No. 17-1387-CFC-SRF, 2018 WL 6617143, at *9 (D. Del. Dec. 17, 2018) (patentee did "not attempt to correlate these alleged improvements [over the prior art] to the claims").[1]

2. The Dependent Claims Include Additional Inventive Concepts

As discussed above, the dependent claims include additional, meaningful limitations that

[1] *Vehicle Intel. & Safety LLC v. Mercedes-Benz USA, LLC*, 635 F. App'x 914, 918 (Fed. Cir. 2015) (patent addressed driver-screening problem "by simply stating 'use an expert system'"); *Quad City Pat., LLC v. Zoosk, Inc.*, 498 F. Supp. 3d 1178, 1185 (N.D. Cal. 2020) (claims directed to the "abstract idea of a service marketplace that uses standardized terms" rather than a specific solution a problem in the field of computers); *Neochloris, Inc. v. Emerson Process Mgmt. LLLP*, 140 F. Supp. 3d 763, 773 (N.D. Ill. 2015) (patentee "provide[d] no explanation or citation as to why these advanced functions are inventive").

are not merely Orca's overgeneralization of "the type or source of the input or output information." D.I. 138 at 18. These include additional inventive concepts, which Orca fails to address at all. D.I. 124, ¶ 108. For example, Orca fails to explain how training LLMs specifically as recited in claims 3, 5, 14, and 16, is routine, conventional, or well-understood.

C. Dismissal Is Not Appropriate

Dismissal is not appropriate, first of all, because the claims are not directed to patent ineligible subject matter as discussed above. In addition, Wiz has, at a minimum, plausibly alleged that the asserted claim elements, alone and in combination, are not well-understood, routine, or conventional to a person of ordinary skill in the art. *See, e.g.*, D.I. 124, ¶¶ 98-99, 103-105. This issue is "a question of fact," and Wiz's allegations "precludes dismissal" under 35 U.S.C. § 101 at the Rule 12(b)(6) stage because the Court is required "to resolve any plausibly alleged factual issues in favor of the patentee." *Blackbird Tech*, 2020 WL 58535 at *6; *Trackthings*, 2023 WL 4926184 at *11(Judgement on the Pleadings is premature when there are disputes of material fact).

The fact that claim construction has not yet occurred is "is another reason why grant of the [Section 101] Motion now would be inadvisable," given the parties' dispute over whether the claims are directed to specific techniques. *See Ficep Corp. v. Peddinghaus Corp.*, C.A. No. 19-1994-RGA, 2021 WL 254104, at *8 (D. Del. Jan. 26, 2021). Orca only asserts in conclusory fashion that claim construction could not impact the Motion without any reference to the language or terms.

VI. CONCLUSION

For the foregoing reasons, Orca's respectfully requests that Orca's Motion be denied in all respects.

Dated: September 19, 2024

Respectfully Submitted

RICHARDS, LAYTON & FINGER, P.A.

/s/ Frederick L. Cottrell, III

Frederick L. Cottrell, III (#2555)
Kelly E. Farnan (#4395)
Christine D. Haynes (#4697)
One Rodney Square
920 N. King Street
Wilmington, DE 19801
(302) 658-6541
cottrell@rlf.com
farnan@rlf.com
haynes@rlf.com

Counsel for Defendant Wiz, Inc.

OF COUNSEL:

Jordan R. Jaffe
Catherine Lacy
Callie Davidson
Alex Miller
WILSON SONSINI GOODRICH & ROSATI, P.C.
One Market Plaza
Spear Tower, Suite 3300
San Francisco, CA 94105
(415) 947-2000
jjaffe@wsgr.com